

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<small>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</small> <b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b>					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
11/26/08		Final Progress Report		68/1/05-7/31/08	
4. TITLE AND SUBTITLE  Stochastic Control and Numerical Methods with Applications to Communications  Game theoretic/subsolution to importance sampling for rare event simulation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER W911NF-05-1-0289	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S)  Harold Kushner and Paul Dupuis				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Brown University Office of Sponsored Projects 164 Angell St. Providence, RI 02912				8. PERFORMING ORGANIZATION REPORT NUMBER  4	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 48657MA	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Game theoretic methods were developed to greatly improve importance sampling for rare event simulation for stochastic systems and queueing networks that arise in communications. With the same goal in mind, large deviations approach to design and analysis was developed to facilitate rare event simulations for systems with discontinuities in the dynamics and analyze escape events for queueing networks. Results for heavy tailed distributions were obtained. Numerical methods were developed for the solution of non-zero sum stochastic game models. Effective stability based methods for scheduling and routing in networks of mobiles with randomly varying channels were produced& they are also useful w/ long term memory systems.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Harold Kushner
					19b. TELEPHONE NUMBER (include area code) 401-863-1400

**Final Report ARO Contract W 911NF-05-1-0289**  
**Harold Kushner and Paul Dupuis**  
**Brown University**  
**November, 2008**

1. **Game theoretic/subsolution approach to importance sampling for rare event simulation.** Importance sampling is a technique that is commonly used to speed up Monte Carlo simulation of rare events. The PIs have established a theoretical framework under which the efficiency of various importance sampling algorithms can be rigorously justified and a flexible constructive methodology by which efficient importance sampling schemes can be built for complex systems. The following papers were completed.
  - (a) *Dynamic importance sampling for queueing networks* (P. Dupuis, D. Sezer, and H. Wang), *Annals of Applied Probability*, **17**, (2007), 1306–1346.
  - (b) *Importance sampling for Jackson networks* (P. Dupuis and H. Wang), to appear in QUESTA.

Papers (a) and (b) are concerned with rare event simulation in the context of queueing networks. This has been an active research area in rare event simulation but little was known regarding the design of efficient importance sampling algorithms for such systems. The standard approach, which simulates the system using an a priori fixed change of measure suggested by large deviation analysis, has been shown to fail in even the simplest network setting (e.g., a two-node tandem network). Exploiting connections between importance sampling, differential games, and classical subsolutions of the corresponding Isaacs equation, we show how to design and analyze simple and efficient dynamic importance sampling schemes for simulating various buffer overflows in stable open Jackson networks. In general, the sampling distributions can be chosen so that they are independent of the particular rare event of interest, and hence overflow probabilities for different events can be estimated simultaneously. A by-product of this type of analysis is the identification of the minimizing trajectory for the calculus of variation problem that is associated with the sample-path large deviation rate function. On the other hand, for systems with special structures such as tandem Jackson networks, importance sampling with better performance can be designed on a case by case basis depending on the particular event under consideration.

2. **Branching methods for fast simulation of rare events.** Besides importance sampling, another important method for simulating rare events is based on branching processes. The PIs have developed a theoretical foundation upon which efficient branching methods can be built for general systems. It was shown that once again subsolutions (albeit not in the

classical sense) to an associated Hamilton-Jacobi-Bellman equation play the key role. Two papers were completed.

- (a) *Splitting for rare event simulation: A large deviations approach to design and analysis* (P. Dupuis and T. Dean), to appear in *Stochastic Processes and their Applications*.
- (b) *The design and analysis of a generalized DPR algorithm for rare event simulation* (P. Dupuis and T. Dean), submitted to *Annals of OR*.

The papers consider branching methods, with or without killing, for rare event simulation. It is assumed that the quantity of interest can be embedded in a sequence whose limit is determined by a large deviation principle. For branching methods, both with killing and without killing, a notion of subsolution is defined for the related calculus of variations problem, and two main results are proved under mild conditions. One is that the number of particles and the total work scales subexponentially in the large deviation parameter when the branching process is constructed according to a subsolution. The second is that the asymptotic performance of the schemes as measured by the variance of the estimate can be characterized in terms of the subsolution. We also compare the methods that use killing with the analogous schemes without killing.

3. **Large deviation and importance sampling for systems with discontinuous dynamics.** Large deviation analysis for systems with discontinuous dynamics in the interior of the state space is difficult. A general theory only exists for systems with two regions of constant statistical behavior separated by a hyperplane of codimension one. The goal of the research in this direction is to, for a large class of physical systems with multiple hyperplanes of discontinuity, give a complete characterization of the large deviation asymptotics. All such systems share a common feature that a type of stability condition that is essential for the large deviation analysis on the discontinuity interfaces is automatically satisfied. Three papers were completed.

- (a) *On the large deviations properties of the weighted-serve-the-longer-queue policy* (P. Dupuis, K. Leder, and H. Wang), to appear in *In and Out of Equilibrium 2*, volume 60 of *Progress in Probability*, Birkhauser, 2008.
- (b) *Importance sampling for weighted-serve-the-longest-queue* (P. Dupuis, K. Leder, and H. Wang) to appear in *Math. of OR*.
- (c) *Large deviations and importance sampling for a tandem network with slow-down* (P. Dupuis, K. Leder, and H. Wang), *QUESTA*, **57**, (2007), 71–83.

Papers (a) and (b) consider a single server system with multi-class arrivals in which the service priority is determined according to the weighted-serve-the-longest-queue policy. The problem setup falls into the general category of systems with discontinuous statistics. Based on a weak convergence approach, we identify the large deviation rate function on the path space. Furthermore, for buffer overflow probabilities of such systems, we explicitly identify the exponential decay rate for the rare-event probabilities of interest, and construct asymptotically optimal importance sampling schemes for simulation based on the game theoretic/subsolution approach.

Paper (c) is concerned with a variant of the two node tandem Jackson network where the upstream server reduces its service rate when the downstream queue exceeds some prespecified threshold. The rare event of interest is the overflow of the downstream queue. Based on a game/subsolution approach, we rigorously identify the exponential decay rate of the rare event probabilities and construct asymptotically optimal importance sampling schemes.

4. **Large deviation analysis for infinite dimensional stochastic systems.** Large deviation analysis for infinite dimensional stochastic systems has often been based on approximations of infinite dimensional partial differential equations, where the technical details can become overwhelming and sometimes unnecessary assumptions are imposed. The PI's approach is based on a variational representation for functionals of Brownian motion. The analysis is technically less demanding and more structured, with less stringent assumptions.

- (a) *Large deviations for infinite dimensional stochastic dynamical systems* (P. Dupuis, A. Budhiraja, and V. Maroulas), *Annals of Probability*, **36**, (2008), 1390–1420.
- (b) *Large deviations for stochastic flows of diffeomorphisms* (P. Dupuis, A. Budhiraja, and V. Maroulas), submitted.

Based on the variational representation, paper (a) studies the large deviation properties for infinite dimensional stochastic differential equations driven by various forms of Brownian noise. Proofs of large deviations properties are reduced to demonstrating basic qualitative properties (existence, uniqueness, and tightness) of certain perturbations of the original process.

Paper (b) establishes a large deviation principle for a general class of stochastic flows in the small noise limit. This result is then applied to a Bayesian formulation of an image matching problem, and an approximate maximum likelihood property is shown for the solution of an optimization problem involving the large deviations rate function.

5. **Large deviation analysis for occupancy models and explicit solutions to a class of related nonlinear partial differential equations.** General occupancy problems, where balls are thrown into urns with various allocation rules, play an important role in many statistical and physical models. The PI is interested in the asymptotic analysis for such systems. Such analysis is important for at least two reasons: to understand the system behavior when the number of balls and the number of urns become large, and to compute quantities of interest using asymptotics. The analysis also reveals a class of nonlinear partial differential equations where explicit or semi-explicit solutions are available. Two papers were completed.

- (a) *Large deviation principle for general occupancy models* (P. Dupuis and J. Zhang), *Combinatorics, Probability and Computing*, **17**, (2008), 437–470.
- (b) *Explicit solutions for a class of nonlinear PDE that arise in allocation problems*, (P. Dupuis and J. Zhang), *SIAM J. on Mathematical Analysis*, **39**, (2008), 1627–1667.

Paper (a) obtains large deviation approximations for the empirical distribution for a general family of occupancy problems including *Maxwell-Boltzmann*, *Bose-Einstein* and *Fermi-Dirac* statistics as special cases. A process level large deviation analysis is conducted and the rate function for the original problem is then characterized, via the contraction principle. This leads to a calculus of variation problem, which is shown to coincide with that of a simple finite dimensional minimization problem. As a consequence, the large deviation approximations and related qualitative information are available in more-or-less explicit form, in sharp contrast to the great majority of large deviation problems for processes with state dependence.

Paper (b) considers the deterministic optimal control problem arising from the large deviation analysis for two classes of allocation problems. The first class considers objects of a single type with a parameterized family of placement probabilities. The second class considers only equally likely placement probabilities but allows for more than one type of object. In both cases, we identify the Hamilton-Jacobi-Bellman equation, whose solution characterizes the minimal cost, explicitly construct solutions, and identify the minimizing trajectories. Paper (b) is also of interest for the reason that it identifies a class of explicitly solvable nonlinear partial differential equations.

6. **Importance sampling for heavy tails.** Random variables with heavy tails differ significantly from those with light tails. For example, the large deviation behavior for heavy tails can be completely different in both scaling and the way the limit optimal paths behave. One paper was completed.

- (a) *Importance sampling for sums of random variables with regularly varying tails* (P. Dupuis, K. Leder, and H. Wang), *ACM Trans. on Modelling and Computer Simulation*, **17**, (2007), 1–21.

The paper is our first attempt at building efficient importance sampling schemes for systems involving heavy-tailed random variables, for which there is little consensus on how to choose the change of measure used in importance sampling. The paper studies state dependent importance sampling schemes for sums of independent and identically distributed random variables with regularly varying tails. The number of summands can be random but must be independent of the summands. For estimating the probability that the sum exceeds a given threshold, we explicitly identify a class of dynamic importance sampling algorithms with bounded relative errors. In fact, these schemes are nearly asymptotically optimal in the sense that the second moment of the corresponding importance sampling estimator can be made as close as desired to the minimal possible value.

7. **Numerical Methods for Non-Zero-Sum Stochastic Differential Games.** In [9, 5] we extended the Markov chain numerical approximation method to non-zero-sum stochastic differential games, where the controls for the two players are separated, a common model. The method is widely used for the numerical solution of stochastic control and optimal control problems in continuous time, for controlled reflected-jump-diffusion type models. It was extended to zero-sum stochastic differential games in [7, 8, 10]. The method has been used in applications of non-zero sum games but it was not known whether it converged to the equilibrium values for the game problem for the original diffusion model. The non-zero sum problem is difficult because each player has its own performance function. The proof for the two-person zero-sum game has the advantage that the controls are determined by a minmax operation and there is a single cost function, so that one player's gain is another's loss, properties that the non-zero-sum game does not have. This difference creates difficulties for the non-zero-sum case that require considerable modification of the proofs. The methods that are employed require the use of strong-sense, rather than with weak-sense solutions, and we must work with strategies and not simply controls.
8. **Numerical Approximations to Optimal Nonlinear Filters.** The usefulness of the theory of nonlinear filtering is limited by the availability of good practical approaches that well approximate the quantities of major interest, for example the conditional (weak-sense) density or the conditional mean and covariance. The mathematical theory is mainly concerned with diffusion-type models and white noise corrupted observations that are taken continuously in time. If the observations are taken in discrete time (as they tend to be in practical applications), then the theoretical issues are less, since one only needs to approximate the (weak-sense) solution to the Fokker-Planck equation between observations and then use Bayes' rule

to incorporate the observations. In [9] we discussed two broad classes of approximations that have been successfully used on various classes of very nonlinear problems and hold considerable promise. The first approach is the so-called Markov chain approximation method based on citeKus-Dup92. At this time it is the most appropriate approach to approximating the weak-sense conditional density, at least for low-dimensional problems. The basic idea is to use a filter for a Markov chain that approximates the diffusion, but with the actual physical observations. Convergence theorems can be proved as the approximation parameters go to zero.

The Markov chain approximation method, when used for the approximation of the conditional weak-sense density, can be computationally intensive. Often one is interested in just the first few conditional moments. Since, for nonlinear problems, a finite set of conditional moments will define the conditional distribution only under very restrictive conditions, one must resort to a heuristic procedure. The second method considered is the so-called “assumed form of the conditional density” approach, first proposed in [6] and developed and used in various ways since then [1, 2, 4, 12]. With this method, one assumes that the conditional density takes a particular parametrized form, and then approximates the evolution of the parameters, under this assumption. Most commonly, the assumed density is Gaussian (or a Gaussian mixture), where the parameters are the conditional mean and covariance. The numerical issues center about the approximation of integrals with respect to Gaussian kernels. With guidance from the literature on the numerical evaluation of integrals, there are many ways of doing this, and some methods of current interest were discussed. Numerical data lend support to the value of the approach.

#### 9. Scheduling and Control of Mobile Communications Networks with Randomly Time Varying Channels by Stability Methods.

Consider a communications network consisting of mobiles, some of which can serve as a receiver and/or transmitter in a multihop path. There are random external data processes, each destined for some destinations. At each mobile the data is queued according to the source-destination pair until transmitted. The capacities of the connecting channels are randomly varying. Time is divided into small scheduling intervals. At the beginning of the intervals, the channels are estimated via pilot signals and this information is used for the scheduling decisions during the interval, concerning the allocation of transmission power and/or time, bandwidth, and perhaps antennas, to the various queues in a queue and channel-state dependent way, to assure stability. General networks are covered, conditions used in previous works were weakened, and the distributions of the input file lengths can be heavy tailed. The resulting controls are readily implementable. The choice of Liapunov function allows a range of tradeoffs between current rates and queue lengths, under very weak conditions.

Owing to the random nature of the arrival and channel processes, the computation or even the existence of stabilizing policies is not at all obvious.

Owing to the non-Markovian nature of the system state, classical stability methods cannot be used without revision, and a perturbed Liapunov function method is adapted to obtain the desired results. This work used a much simpler Liapunov function perturbation, based on a simple mixing condition, that has many advantages and allows us to deal with processes not covered by previous work. It is more manageable, and it extends the methods so that heavy tailed input processes can be handled.

With this method, and  $X$  denoting the vector of queue values at all the nodes, one starts with a basic Liapunov function  $V(X)$  that works for a “mean flow” system. Then one gets a perturbation  $\delta V(n)$  to  $V(X)$  so that  $V(X(n)) + \delta V(n)$  can be used as a Liapunov function for the actual non-Markov physical system and imply the desired stability. The actual decision rule is based on the gradient of  $V(X)$  and is readily implemented. The basic result is that, if a certain “mean flow” or fluid approximation process is stable, then so is the physical system under our scheduling rule. This stabilizability of the mean flow approximation can often be readily verified. The condition is nearly necessary as well.

10. **Numerical methods for optimal controls for stochastic systems with delays.** This was a major part of our effort, and led to the comprehensive book [3]. It is an extension to the model with delays of the Markov chain approximation methods of [11]. For the nondelay problem, these methods are a widely used and powerful class of numerical approximations of optimal costs or other functionals of controlled or uncontrolled stochastic processes in continuous time. There are numerous sources of delays in the modeling of realistic physical and biological systems. Many examples arise in communications and queueing, due to the finite speed of signal transmission, the nonnegligible time required to traverse long communications distances, or the time required to go through a queue. Other examples arise because of mechanical transportation delays as, for example in hydraulic control systems, delays due to noninstantaneous human responses or chemical reactions, or delays due to visco-elastic effects in materials. Very little information is available concerning solutions when the models are nonlinear and stochastic, and numerical methods should be a main source of such information.

There is a huge literature on control problems for delay systems for the linear model (deterministic or stochastic) with a quadratic cost criterion, and many good computational methods have been developed. Although these techniques and algorithms have been useful for the linear problem, it is not clear how to adapt them to the nonlinear models that are of concern to us. For this reason, we confined attention to analogs of the approaches that have been found to be very useful for the general no-delay problem, namely the Markov chain approximation method.

The models of the systems of concern are diffusion and reflected diffusion processes, and the results can be extended to cover jump-diffusions. The



control might be “ordinary” in the sense that it is a bounded measurable function, or it might be impulsive, or what is known as a “singular” control. All of the usual cost functionals are covered; the discounted cost, stopping on reaching a boundary, optimal stopping, ergodic, etc. Any or all of the path, control, boundary reflection process, or driving Wiener process, might appear in delayed form. Examples where the boundary reflection process might be delayed occur in communications/queueing models, where there is a communications delay. If a buffer overflows (corresponding to a lost packet), a signal is sent to the source, which receives it after a delay, and then adjusts its rate of transmission accordingly. The buffer overflow is a component of the boundary reflection process. Models with delays of such boundary reflection terms have not been treated previously.

For the nondelay problem, the approach of the Markov chain approximation method starts by approximating the original controlled process by a controlled Markov chain on a finite state space. The approximation parameter is denoted by  $h$  and it might be vector-valued. The original cost functional is also approximated so that it is suitable for the chain. The approximating chain must satisfy a simple condition called “local consistency.” This is quite unrestrictive and means simply that from a local point of view and for small  $h$ , the conditional mean and covariance of the changes in state of the chain are proportional to the local mean drift and covariance of the original process, modulo small errors. Many straightforward ways of getting the approximating chains are discussed in [11], where it is seen that the approach is very flexible. The approximation yields a control problem that is close to the original, which gives the method intuitive content that can be exploited for the construction of effective algorithms. After getting the approximating chain, one solves the Bellman equation for the optimal cost (or simply the equation for the value function of interest if there is no control), and proves that the solution converges to the desired optimal cost or value function as  $h$  goes to zero. One tries to choose the approximation so that the associated control or optimal control problem can be solved with a reasonable amount of computation and that the approximation errors are acceptable.

The proofs of convergence of the Markov chain approximation method as  $h \rightarrow 0$  are purely probabilistic. We always work with the processes. No tools from PDE theory or classical numerical analysis are used. The idea behind the proof can be described as follows. For the optimal control problem, starting with the approximating chain with its optimal control, one gets a suitable continuous-time interpolation, and shows that in the sense of weak or distributional convergence, there is a convergent subsequence whose limit is an optimally controlled process of the original diffusion type, and with the original cost function and boundary data. The mathematical basis is the theory of weak convergence of probability measures, and this powerful theory provides a unifying approach for all of the problems of

interest.

The probabilistic nature of the methods of process approximation and of the mathematical proofs of convergence allows us to use our physical intuition concerning the original problem in all phases of the development. This gives us great flexibility in the details of the approximation and in the construction of algorithms. These advantages will carry over to the problem with delays. In fact, the probabilistic approach to the approximation and convergence is particularly important when there are delays, since virtually nothing is known about the analytical properties of the associated (infinite-dimensional) Bellman equations for nonlinear problems.

For models without delays, the system state takes values in a subset of some finite-dimensional Euclidean space, and the control is a functional of the current state. For models with delays, the state space must take the path of the delayed quantities (over the delay intervals) into account, and this makes the problem infinite-dimensional. So a major issue in adapting the Markov chain approximation method to models with delays concerns suitable “finite” approximations to the “memory segments” so that a reasonable numerical method can be devised, and much attention is given to this problem. The methods of approximation that were developed are natural and seem to be quite promising. They deal with issues of approximation that are fundamental.

Suppose that the effect of the control action is delayed. This can cause serious instabilities. To effectively control in such a case, in determining the current control action one must take into account the control actions that were made in the recent past but whose effects have not yet been seen by the controller, those up to the maximum delay interval back from the present time.

The book summarizes the main results that will be needed from the theory of weak convergence of a sequence of random processes. The primary processes of concern in the proofs of convergence are continuous-time interpolations of the approximating chains, and we will need to show that they have limits that are (in fact, optimal) controlled diffusions. Weak convergence theory, together with the methods of the so-called martingale problem for characterizing the limit process as the desired diffusions, provides the essential tools. With their use, the proofs of convergence are purely probabilistic. For the no-delay case this probabilistic approach to the proofs of convergence of numerical algorithms is the most powerful and flexible. For the delay case, there does not seem to be any alternative since the Bellman equation is infinite-dimensional and virtually nothing is known about it.

The existence of an optimal control is also shown. The proof of this fact is important because it is a template for the proofs of convergence of the system and numerical approximations in subsequent chapters. For the singular control problem, the definition of the model and the existence

of an optimal control are dealt with via a very useful “time transformation” method, which is necessary owing to the possibly wild nature of the associated paths and controls.

The key difference between the problem with and without delays is that the state space for the problem with delays involves the “memory segments” of the components whose delayed values appear in the dynamics. The first step in the construction of a numerical approximation involves approximating the original dynamical system. In our case, this entails approximating the delays and dynamics so that the resulting model is simpler, and ultimately finite-dimensional. We develop model simplifications that have considerable promise when the path or path and/or control are delayed. A variety of approximations are presented, eventually leading to finite-dimensional forms that are used as the basis of numerical algorithms. To help validate the approximations, simulations that compare the paths of the original and approximated system were presented, and it is seen that the approximations can be quite good.

Delay equations might have rapidly time-varying terms, even rapidly varying delays. This complicates the numerical problem. But, under suitable conditions, there are limit and approximation theorems that allow us to replace the system by a simpler “averaged” one and some such results are presented.

The average cost per unit time (ergodic cost) problem for nondegenerate reflected diffusion models, where only the path is delayed, is developed. There are only a few results on the ergodic theory for general delay equations. Since they are not adequate for the needs of the numerical and approximation problems for the systems of interest, the necessary results are developed, using methods based on the Girsanov transformation and the Doeblin condition. Of particular interest is the demonstration that the various model approximations developed for the non-ergodic problem can also be used.

Owing to the local consistency condition, the dynamical system that is represented by a continuous-time interpolation of the chain “resembles” the original controlled diffusion process. Thus we would expect that the optimal cost or the values of the functionals of interest would be close to those for the diffusion. This is quantified by the convergence theorems. There are two (asymptotically equivalent) methods of getting the approximating chains that are of interest, called the “explicit” and “implicit” methods. They differ in the way that the time variable is treated, and each can be obtained from the other. The first method was the basic approach for the nondelay problem. The second method plays a useful role in reducing the memory requirements when there are delays.

It is shown that any method of constructing the approximating chain for the no-delay problem can be readily adapted to the delay problem, with the transition probabilities taking the delays into account. The only

change in the local consistency condition is the use of the “memory segment” arguments in the drift and diffusion functions. The algorithms are well motivated and seem to be quite reasonable. But since the subject is in its infancy, what was presented should be taken as a first step, and will hopefully motivate further work. When constructing a numerical approximation algorithm, there are two main issues that must be kept in mind. The algorithm must be numerically feasible and it must be such that there is a proof of convergence as the approximating parameter goes to zero. These issues inform the structure of the development. A large variety of numerical approximations are developed, always keeping an eye on the memory size problem.

The memory requirements can become onerous if the reflection process and/or the Wiener process also appear in delayed form, or if the control-value space has more than a few points. We developed an alternative approach that reduces the memory requirements for general nonlinear stochastic problems where the control and reflection terms, as well as the path variables, are delayed. Effectively, the delay equation is replaced by a type of stochastic wave equation with no delays, and its numerical solution yields the optimal costs and controls for the original model. The representation is equivalent to the original problem in that any solution to one yields a solution to the other. The details of the appropriate Markov chain approximation are given and the convergence theorem is proved. Theoretically, with the use of appropriate numerical approximations, the dimension of the required memory vector is much reduced.

## References

- [1] I. Arasaratnam, S. Haykin, and R. J. Elliott. Discrete-time nonlinear filtering algorithms using Gauss-Hermite quadrature. *IEEE Transactions*, 95:953–977, 2007.
- [2] A. Budhiraja and H.J. Kushner. A nonlinear filtering algorithm based on an approximation of the conditional distribution. *IEEE Trans. Automatic Control*, 45:580–585, 2000.
- [3] H.J. Kushner. *Numerical Methods for Controlled Stochastic Delay Systems*. Birkhäuser, Boston, 2008.
- [4] K. Ito and K. Xiong. New gaussian filters for nonlinear filtering problems. Submitted to *IEEE Trans. Automatic Control*, 1997.
- [5] H. J. Kushner. Numerical approximations for non-zero-sum stochastic differential games. *SIAM J. on Control and Optimiz.*, 40:1942–1971, 2007.
- [6] H.J. Kushner. Approximations to optimal nonlinear filters. *IEEE Trans. Automatic Control*, AC-12:546–556, 1967.

- [7] H.J. Kushner. Numerical approximations for stochastic differential games. *SIAM J. Control Optim.*, pages 457–486, 2002.
- [8] H.J. Kushner. Numerical approximations for stochastic differential games: The ergodic case. *SIAM J. Control Optim.*, 42:1911–1933, 2004.
- [9] H.J. Kushner. Numerical methods for non-zero-sum stochastic differential games: Convergence of the Markov chain method. In *Topics in stochastic analysis and nonparametric estimation*. Springer-Verlag, Berlin and New York, 2007. An IMA volume, to appear.
- [10] H.J. Kushner. Numerical methods for stochastic differential games: The ergodic cost criterion. In S. Jorgensen, M. Quincampoix, and T. Vincent, editors, *Advances in Dynamic Game Theory: Numerical Methods, Algorithms, and Applications to Ecology and economics*, pages 617–638. Birkhauser, Boston, 2007.
- [11] H.J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, Berlin and New York, 1992. Second edition, 2001.
- [12] H. Tanizaki. *Nonlinear Filters: Estimation and Applications*. Springer Verlag, Berlin and new York, 1996.